

Implementing SDS Tips and Tricks



Jeff Filippi
Integrated Data Consulting, LLC

Session: A05
Monday, April 26, 2010: 3:30 – 4:30

Implementing SDS Tips and Tricks

Session: A05

Jeff Filippi



Integrated Data Consulting, LLC
jeff.filippi@itdataconsulting.com

Introduction

- 20 years of working with Informix products
- 15 years as an Informix DBA
- Worked for Informix for 5 years 1996 – 2001
- Certified Informix DBA
- Started my own company in 2001 specializing in Informix Database Administration consulting services
- IBM Business Partner
- OLTP and Data warehouse systems
- Informix 4.x, 5.x., 7.x, 9.x, 10.x, 11.10, 11.50



Overview

- Why do you need to implement SDS
- Potential changes needed for your application
- Monitoring SDS environments
- Issues I have come across in SDS implementations



Why do you need to implement SDS

- Need to increase capacity of your business and have the ability to add (n) number of instances.
- Do not want to duplicate the disks for another system.
- Do not want to change your schema to have to have primary keys like in ER.
- You want to run different types of processes on each instance.

Why do you need to implement SDS – (cont'd)

- Here are a couple scenarios on how to use SDS:
 - Scenario #1 (Separate data loading/batch from OLTP)
 - Primary – Data loading/batch jobs
 - SDS – User queries/Reporting/OLTP type operations
 - Scenario #2 (Balance load between servers)
 - Primary – OLTP users
 - SDS – OLTP users
 - Thru the use of the connection manager, you can manage balancing the load between the servers.



Potential changes needed for your Application

- Check for error code -7350 in the application for the case when the before image on the secondary is different than the current image on the primary and the write operation is not allowed.
- Add VERCOLS to the schema of the tables. This allows only a changed column to be compared and not the whole row. This reduces network traffic.
- Check to see if you application uses home grown sequence generators. These may need to be changed.
- Does your application create/drop tables or indexes?
 - If so you would need to change your application if you wanted it to run on an SDS instance since it does not allow DDL operations.

Potential changes needed for your Application – (cont'd)

- Error Code -7350
 - Error Description:
 - Attempt to update a stale version of a row
 - An attempt was made to update a stale copy of a row. This caused an optimistic concurrency failure.
 - This error can occur when using updatable secondary and the current version of the row has not yet been replicated to the secondary on which the client application is connected.
 - This error code is also returned when table schema at secondary server doesn't match with the table schema at primary server.
- Work Around
 - Change application to retry the SQL.
 - In the case of my customer they put a check in their code that if they received the error -7350 they would retry the SQL statement up to 5 times.



Potential changes needed for your Application – (cont'd)

- Issue with home grown sequence generators.
 - I had a customer who we implemented SDS with and ran into a problem with their own sequence number generator.
 - After go live, the customer noticed in about a ½ % of the time that they were seeing duplicates.
 - Since SDS uses optimistic locking, the table that was keeping the counter had multiple processes updating the table and since the record was not locked during the update it was being reset.
 - Example:
 - On SDS - Process 1: select seq_cntr from seq_cntr_tbl where cntr_id = 4 - (Returns 1500)
 - On SDS - Process 2: select seq_cntr from seq_cntr_tbl where cntr_id = 4 - (Returns 1500)
 - Process 1 updates the record to 1501.
 - Process 2 updates the record to 1501.



Potential changes needed for your Application – (cont'd)

- Resolution to home grown sequence generators.
- Use sequence database objects:
 - Replaced the select with the sequence object:
 - Create sequence ser_num increment by 1 maxvalue 9223372036854775807 min value 1913394 cache 20 order;
 - When inserting into a table use the following syntax:
 - Insert into test1 (bat_num, name) values (ser_num.NEXTVAL, "test")

Instead of:

- select seq_cntr from seq_cntr_tbl where cntr_id = 4
- Update seq_cntr_tbl set seq_cntr = seq_cntr + 1 where cntr_id = 4
- Insert into test1 (bat_num, name) values (seq_cntr, "test")



Potential changes needed for your Application – (cont'd)

- Adding VERCOLS to tables
 - Alter table xyz add vercols;
 - Create table xyz (col1 integer, col2 integer) with vercols;



Monitoring SDS Environments

- Options in monitoring SDS Environments
 - You can monitor SDS environments thru:
 - onstat
 - sysmaster tables
 - OAT
 - Use of SDS_PAGING and SDS_TEMPDBS



Monitoring SDS Environments – (cont'd)

- ONSTAT
 - On Primary
 - `onstat -g sds`
 - `onstat -g sds verbose`
 - On SDS
 - `onstat -g sds`
 - `onstat -g sds verbose`



Monitor SDS from Primary

- `onstat -g sds`

IBM Informix Dynamic Server Version 11.50.UC6DE -- On-Line (Prim) -- Up
06:43:23 -- 314888 Kbytes

Local server type: Primary
Number of SDS servers:1

SDS server information

SDS srv name	SDS srv status	Connection status	Last LPG sent (log id,page)	Supports Proxy Writes
sds1	Active	Connected	416,9395	Y



Monitor SDS from Primary

- `onstat -g sds verbose`

IBM Informix Dynamic Server Version 11.50.UC6DE -- On-Line (Prim) -- Up 06:43:28 -- 314888
Kbytes

Number of SDS servers:1

Updater node alias name :production

SDS server control block: 0x467120b8

server name: sds1

server type: SDS

server status: Active

connection status: Connected

Last log page sent(log id,page):416,9395

Last log page flushed(log id,page):416,9395

Last log page acked (log id, page):416,9395

Last LSN acked (log id,pos):416,38482196

Approximate Log Page Backlog:0

Current SDS Cycle:1

Acked SDS Cycle:1

Sequence number of next buffer to send: 32

Sequence number of last buffer acked: 31

Time of last ack:2010/02/11 15:02:40

Supports Proxy Writes: Y



Monitor SDS from SDS

- `onstat -g sds`

IBM Informix Dynamic Server Version 11.50.UC6DE -- Updatable
(SDS) -- Up 00:02:59 -- 73588 Kbytes

Local server type: SDS

Server Status : Active

Source server name: production

Connection status: Connected

Last log page received(log id,page): 416,9395



Monitor SDS from SDS

- `onstat -g sds verbose`

IBM Informix Dynamic Server Version 11.50.UC6DE -- Updatable (SDS) --
Up 00:03:04 -- 73588 Kbytes

SDS server control block: 0x462a4c98

Local server type: SDS

Server Status : Active

Source server name: production

Connection status: Connected

Last log page received(log id,page): 416,9395

Next log page to read(log id,page):416,9396

Last LSN acked (log id,pos):416,38482196

Sequence number of last buffer received: 37

Sequence number of last buffer acked: 37

Current paging file:./sds1_paging_1

Current paging file size:8192

Old paging file:./sds1_paging_2

Old paging file size:0



Monitoring SDS Environments – (cont'd)

- SYSMASTER
 - Primary - syssrcsds

address	1181356216
server_name	sds1
server_status	Active
connection_status	Connected
last_sent_log_uniq	417
last_sent_logpage	2275
last_flushed_log_uniq	417
last_flushed_logpage	2275
last_acked_lsn_uniq	417
last_acked_lsn_pos	9318424
seq_tosend	38
last_seq_acked	37
timeof_lastack	1266029329
totallsn_posted	0
totallsn_sent	0
totalpageflush_posted	0
totalpageflush_sent	0

Monitoring SDS Environments – (cont'd)

- SYSMASTER
 - SDS - systrgsds

address	1177177240
source_server	production
connection_status	Connected
last_received_log_uniq	417
last_received_log_page	2278
next_lpgtoread_log_uniq	417
next_lpgtoread_log_page	2279
last_acked_lsn_uniq	417
last_acked_lsn_pos	9330756
last_seq_received	112
last_seq_acked	112
cur_pagingfile	./sds1_paging_2
cur_pagingfile_size	14336
old_pagingfile	./sds1_paging_1
old_pagingfile_size	0

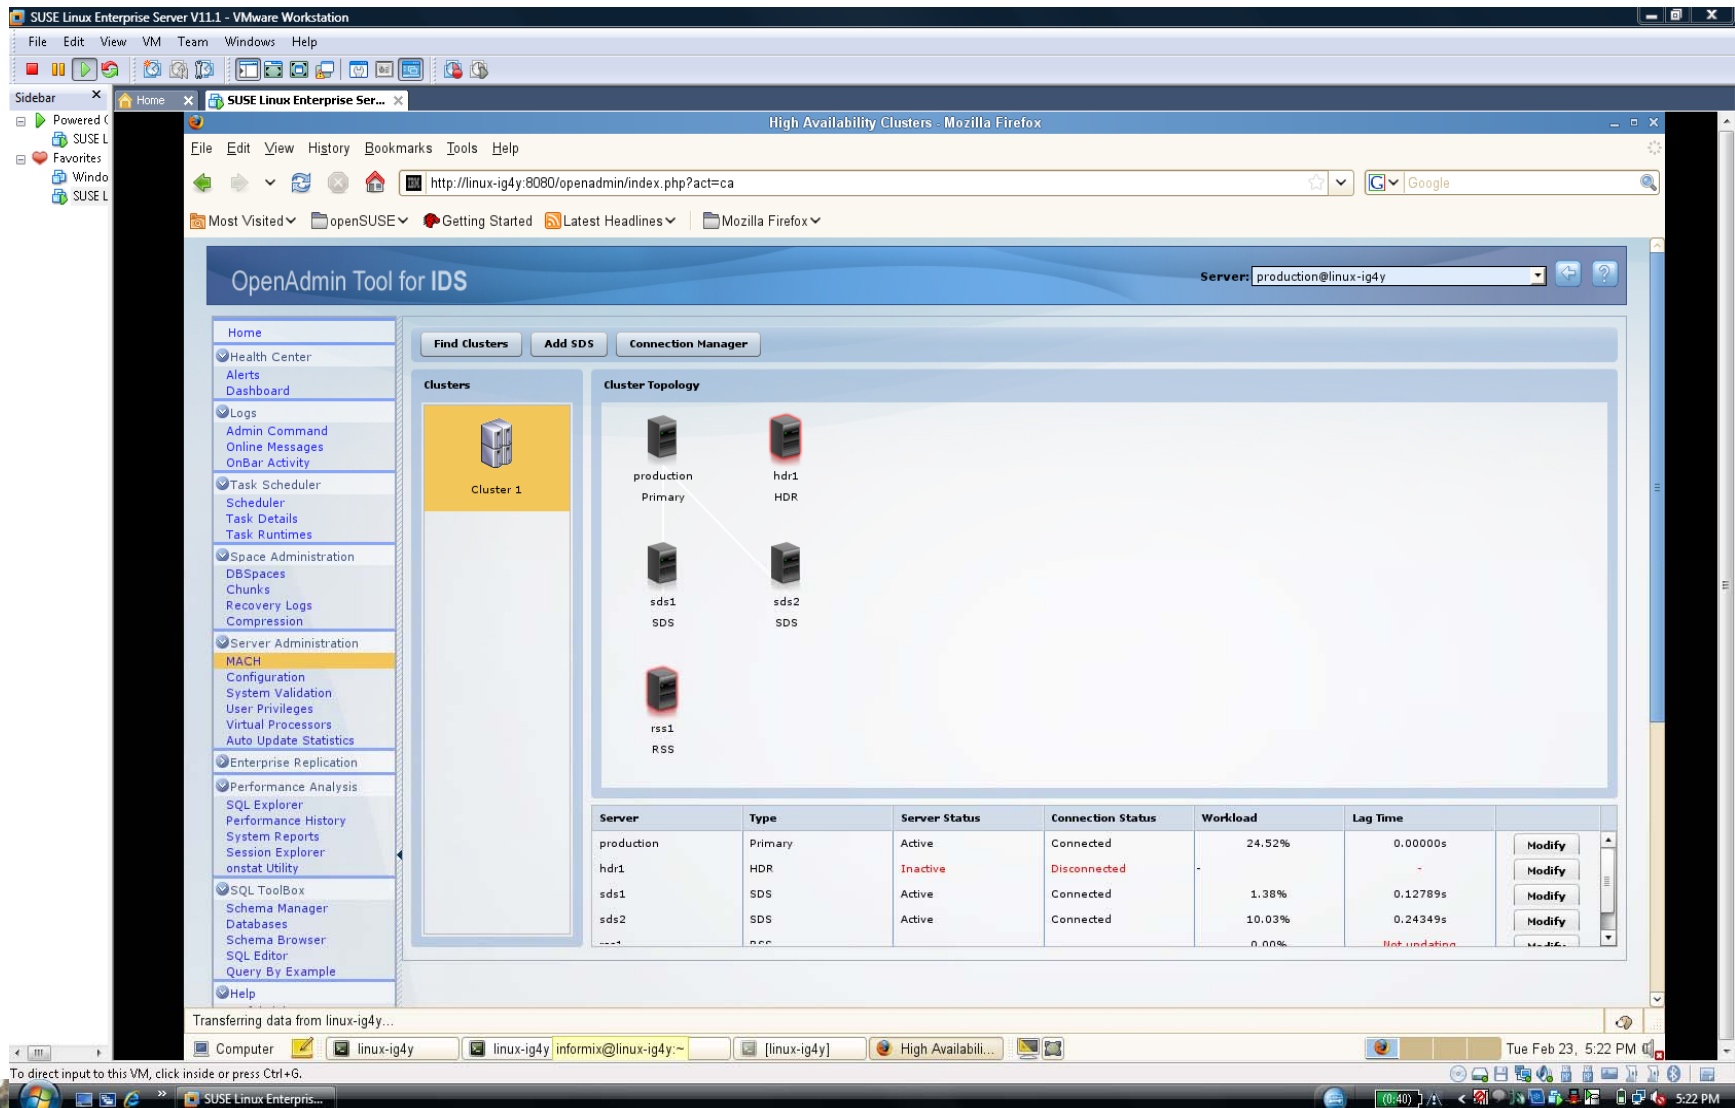


Monitoring SDS Environments – (cont'd)

- Recovery Threads
 - OFF_RECVRY_THREADS:
 - These are the threads which apply the log records on the secondary. If care is not taken, then there will be an imbalance on usage by the recovery threads which will result in a bottleneck. Care must be taken to reduce an imbalance in recovery thread usage or you will end up underutilizing the resources on the secondary. This can result in a backflow which will impact the primary.
 - **Make OFF_RECVRY_THREADS at least 3 times the number of CPUVPS (minimum of 11)**
 - These are the threads which apply the log records on the secondary. Don't be skimpy or apply performance will suffer.
 - Use `onstat -g cpu` to check if one of the recovery threads is doing the bulk of the work.
 - **Consider using a near prime number of recovery threads (i.e. not divisible by 2,3,5,7,11)**
 - This can help minimize sine-wave usage patterns from developing.



Monitor SDS thru OAT



The screenshot displays the OpenAdmin Tool for IDS interface within a Mozilla Firefox browser window. The interface shows a cluster topology diagram and a table of server status.

Cluster Topology:

- Cluster 1 (Primary)
- hdr1 (HDR)
- sds1 (SDS)
- sds2 (SDS)
- rss1 (RSS)

Server Status Table:

Server	Type	Server Status	Connection Status	Workload	Lag Time	Actions
production	Primary	Active	Connected	24.52%	0.00000s	Modify
hdr1	HDR	Inactive	Disconnected	-	-	Modify
sds1	SDS	Active	Connected	1.38%	0.12789s	Modify
sds2	SDS	Active	Connected	10.03%	0.24349s	Modify
...	0.00%	Not updating	...

Monitoring SDS Environments (cont'd)

- Verify sizing for SDS_TEMPDBS, SDS_PAGING
 - **SDS_PAGING** - specifies the path to two files that are used to hold pages that may need to be flushed between checkpoints. Each file acts as temporary disk storage for chunks of any page size.
 - **SDS_TEMPDBS** - The temporary dbspaces are created (or initialized if the dbspaces existed previously) when the SD secondary server starts. The temporary dbspaces are used for creating temporary tables. There must be at least one occurrence of the SDS_TEMPDBS configuration parameter in the ONCONFIG file of the SD secondary server for the SD secondary server to start. You can specify up to 16 SD secondary dbspaces in the ONCONFIG file by using multiple occurrences of the SDS_TEMPDBS configuration parameter.
- Also set these up on the primary server also in case it would become a SDS instance.



Issues I Have Come Across

- Finding out what product to use to share the disks.
 - Last year when I was planning the implementation of SDS I ran into the issue that there was not much documentation on what products were needed to share the disks.
 - In my case the client was using HP-UX Itanium 11.23.
 - In order to share RAW disks the only product that I was able to find was “HP Serviceguard Extention for RAC”. **NOTE:** Also figure this into your total cost of implementing SDS, it is not cheap.
 - There was a presentation “IDS Availability Best Practices” by Madison Pruet on November 18th that did give some more details on this.



Issues I Have Come Across – (cont'd)

- Outages needed for SAN configuration
 - With the implementation of using the clustering software in our case HP Serviceguard Extention for RAC, when the UNIX admin's needed to add new disks/raw devices in the cluster environment, they would have to shutdown the cluster and re-import the volume groups which required an outage of ALL the instances (Primary/SDS/HDR).



Issues I Have Come Across – (cont'd)

- Performance Issue
 - One client had implemented 2 SDS instances and was having performance issues on one of the SDS instances by not the other. The onconfig was the same, the hardware was the same, so what was the issue?
 - They had configured the network at 100/mbs on the one that was slow and 1000/mbs on the fast one. OOPS!!



Issues I Have Come Across – (cont'd)

- On SDS instance a client was seeing timeouts in the application about once a day.
 - After researching the issue I noticed that the time they were experiencing this was when a checkpoint was occurring.
 - In fact for this instance that was the only time of day a checkpoint was occurring.
 - Had `RTO_SERVER_RESTART` set to 1800
 - The logical logs were 10 gig in size and would fill in approximately 24 hour period. When the checkpoint occurred, it was when all the logs but a couple had filled since the last checkpoint. The checkpoint was blocking transactions to make sure that the logical logs did not wrap around since the last checkpoint.



Issues I Have Come Across – (cont'd)

– RESOLUTION:

- In order to fix the issue, we first tried to force a checkpoint (onmode -c) every hour and the issue went away.
- Also to force a checkpoint more often would be to reduce the RTO_SERVER_RESTART, in our case we reduced it to 300 which also had the instance checkpointing more often throughout the day and we did not see the issue again.



Issues I Have Come Across – (cont'd)

- Setting of UPDATABLE_SECONDARY on SDS Instances
 - I had tried to set UPDATABLE_SECONDARY to 2 times the number of CPUVP's but I had an issue where it caused performance degradation.
 - After a discussion with Madison Pruet, I changed it to 4 and the issue went away.



Issues I Have Come Across – (cont'd)

- You can create a RAW table
- BUT – You Cannot change a RAW table to STANDARD
- You need to shutdown the SDS instances.
 - -19845: You cannot alter the logging mode of a table in a logged database on a primary server.



Issues I Have Come Across – (cont'd)

- Using HPL with Express Mode
 - I came across an issue where my HPL job was failing when I was trying to load data into a table.
 - After opening a case with Tech Support they were able to determine the issue was due to having “vercols” on the table.
 - RESOLUTION:
 - Add the hidden columns “ifx_insert_checksum” and “ifx_row_version” to the select statement.
 - BEGIN OBJECT QUERY hist_tbl
 - PROJECT hist_tbl_1
 - DATABASE tst
 - SELECTSTATEMENT "select *, ifx_insert_checksum, ifx_row_version from hist_tbl"



Summary

So in conclusion:

- Why do you need to implement SDS
- Potential changes needed for your application
- Monitoring SDS environments
- Issues I have come across in SDS implementations



QUESTIONS??



Implementing SDS Tips and Tricks

Session: A05

Jeff Filippi



Integrated Data Consulting, LLC

jeff.filippi@itdataconsulting.com

www.itdataconsulting.com

630-842-3608

